Artificial Intelligence & Image Creation

George Legrady © 2020
Experimental Visualization Lab
Media Arts & Technology
University of California, Santa Barbara

Nov 12 **Introduction to the Topic**
Weihao will present on Convolutional Neural-Networks (CNN)

Nov 17 **Review of Current Aesthetic Directions**

Nov 19 **Fabian Offert** lecture on AI & Humanities Perspective
Mert Toka, Neural Texture Synthesis overview

Nov 24 **Review of DeepFakes & Social Implications**

Digital image *made up of pixels* is a multi-dimensional data structure

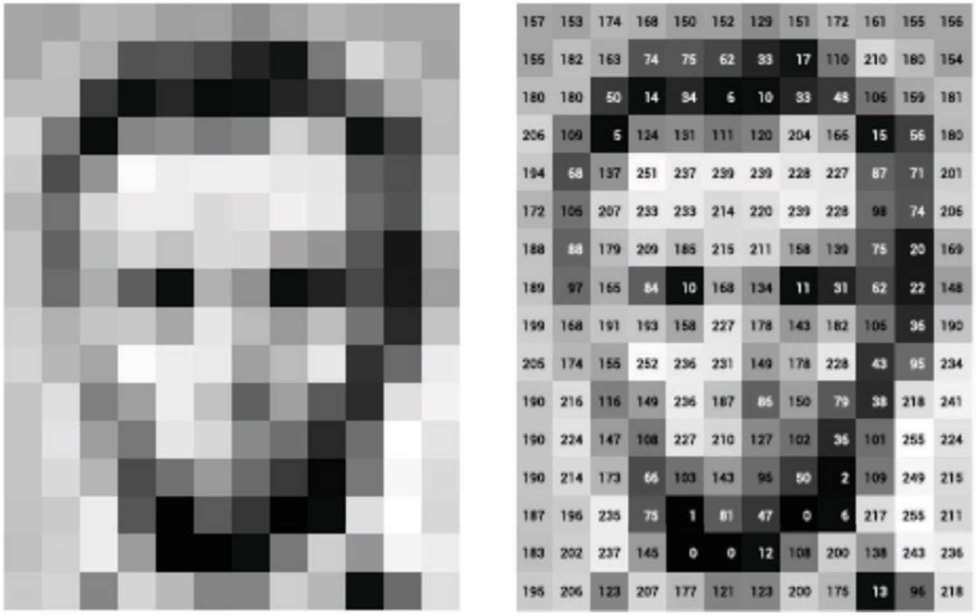

157	153	174	168	150	152	129	151	172	161	155	156
155	182	163	74	75	62	33	17	110	210	180	154
180	180	50	14	34	6	10	33	48	106	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	106	36	190
205	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	86	150	79	38	218	241
190	224	147	108	227	210	127	102	36	101	255	224
190	214	173	66	103	143	96	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	200	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

- Pixel ***Horizontal*** location
- Pixel ***Vertical*** location
- Pixel ***Red*** color value
- Pixel ***Green*** color value
- Pixel ***Blue*** color value
- Pixel ***Alpha*** (transparency) value
- The whole image has a ***BitDepth*** resolution (2bit, 16bit, etc.)

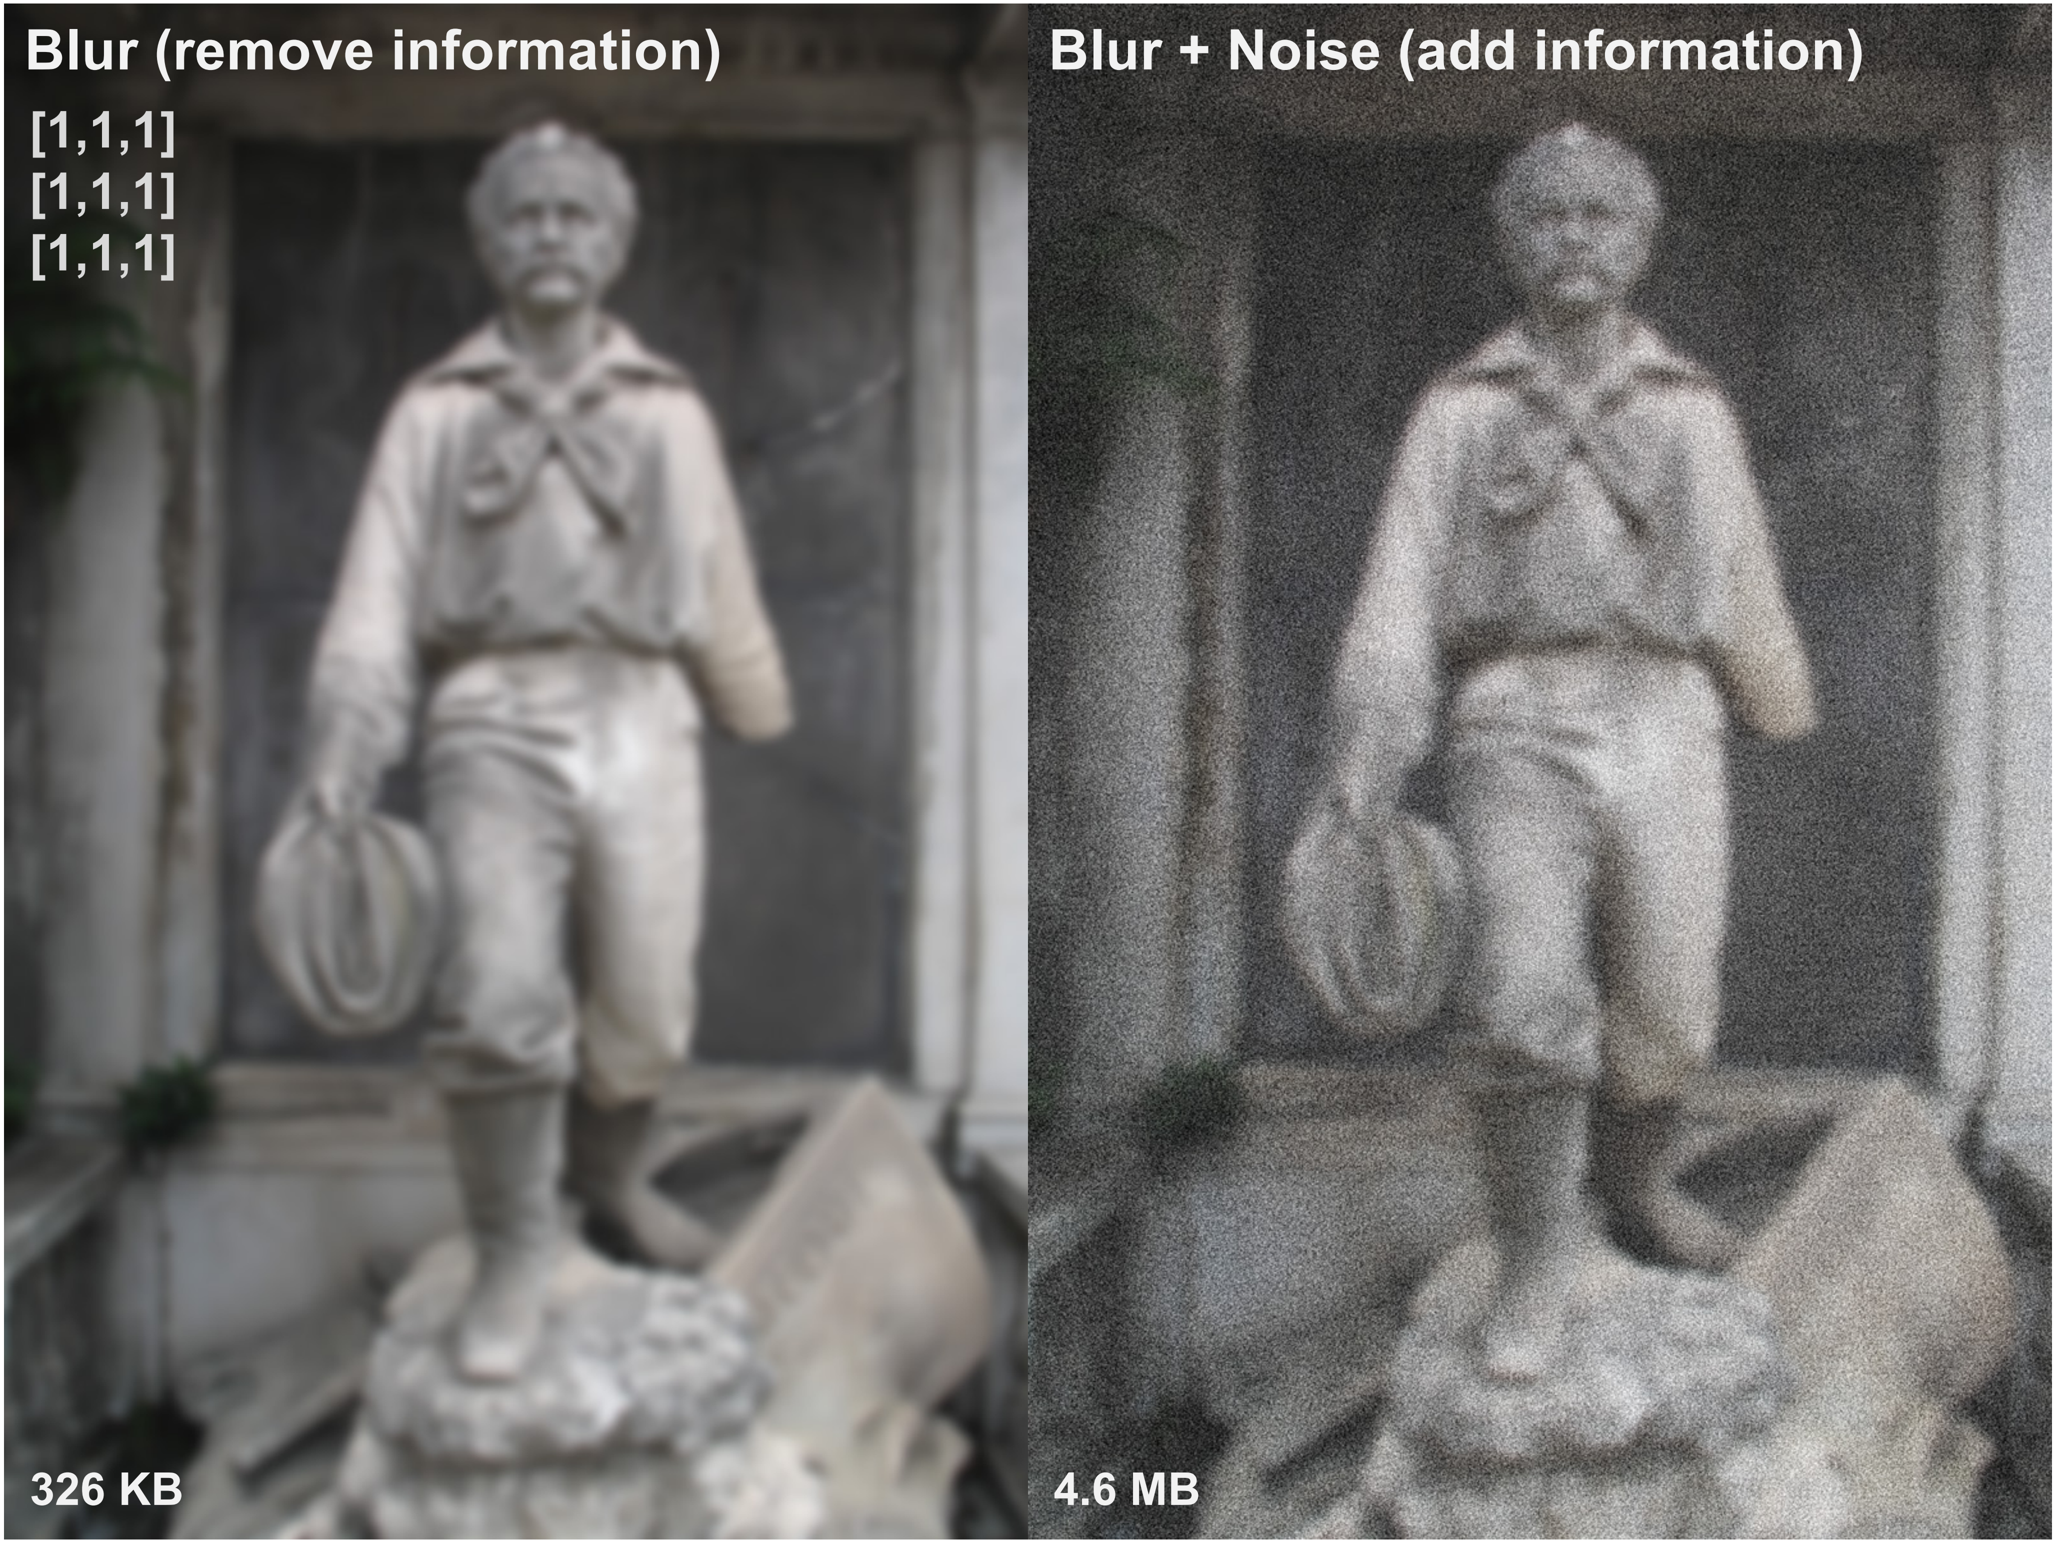

Blur (remove information)
[1,1,1]
[1,1,1]
[1,1,1]
326 KB
Blur + Noise (add information)
4.6 MB

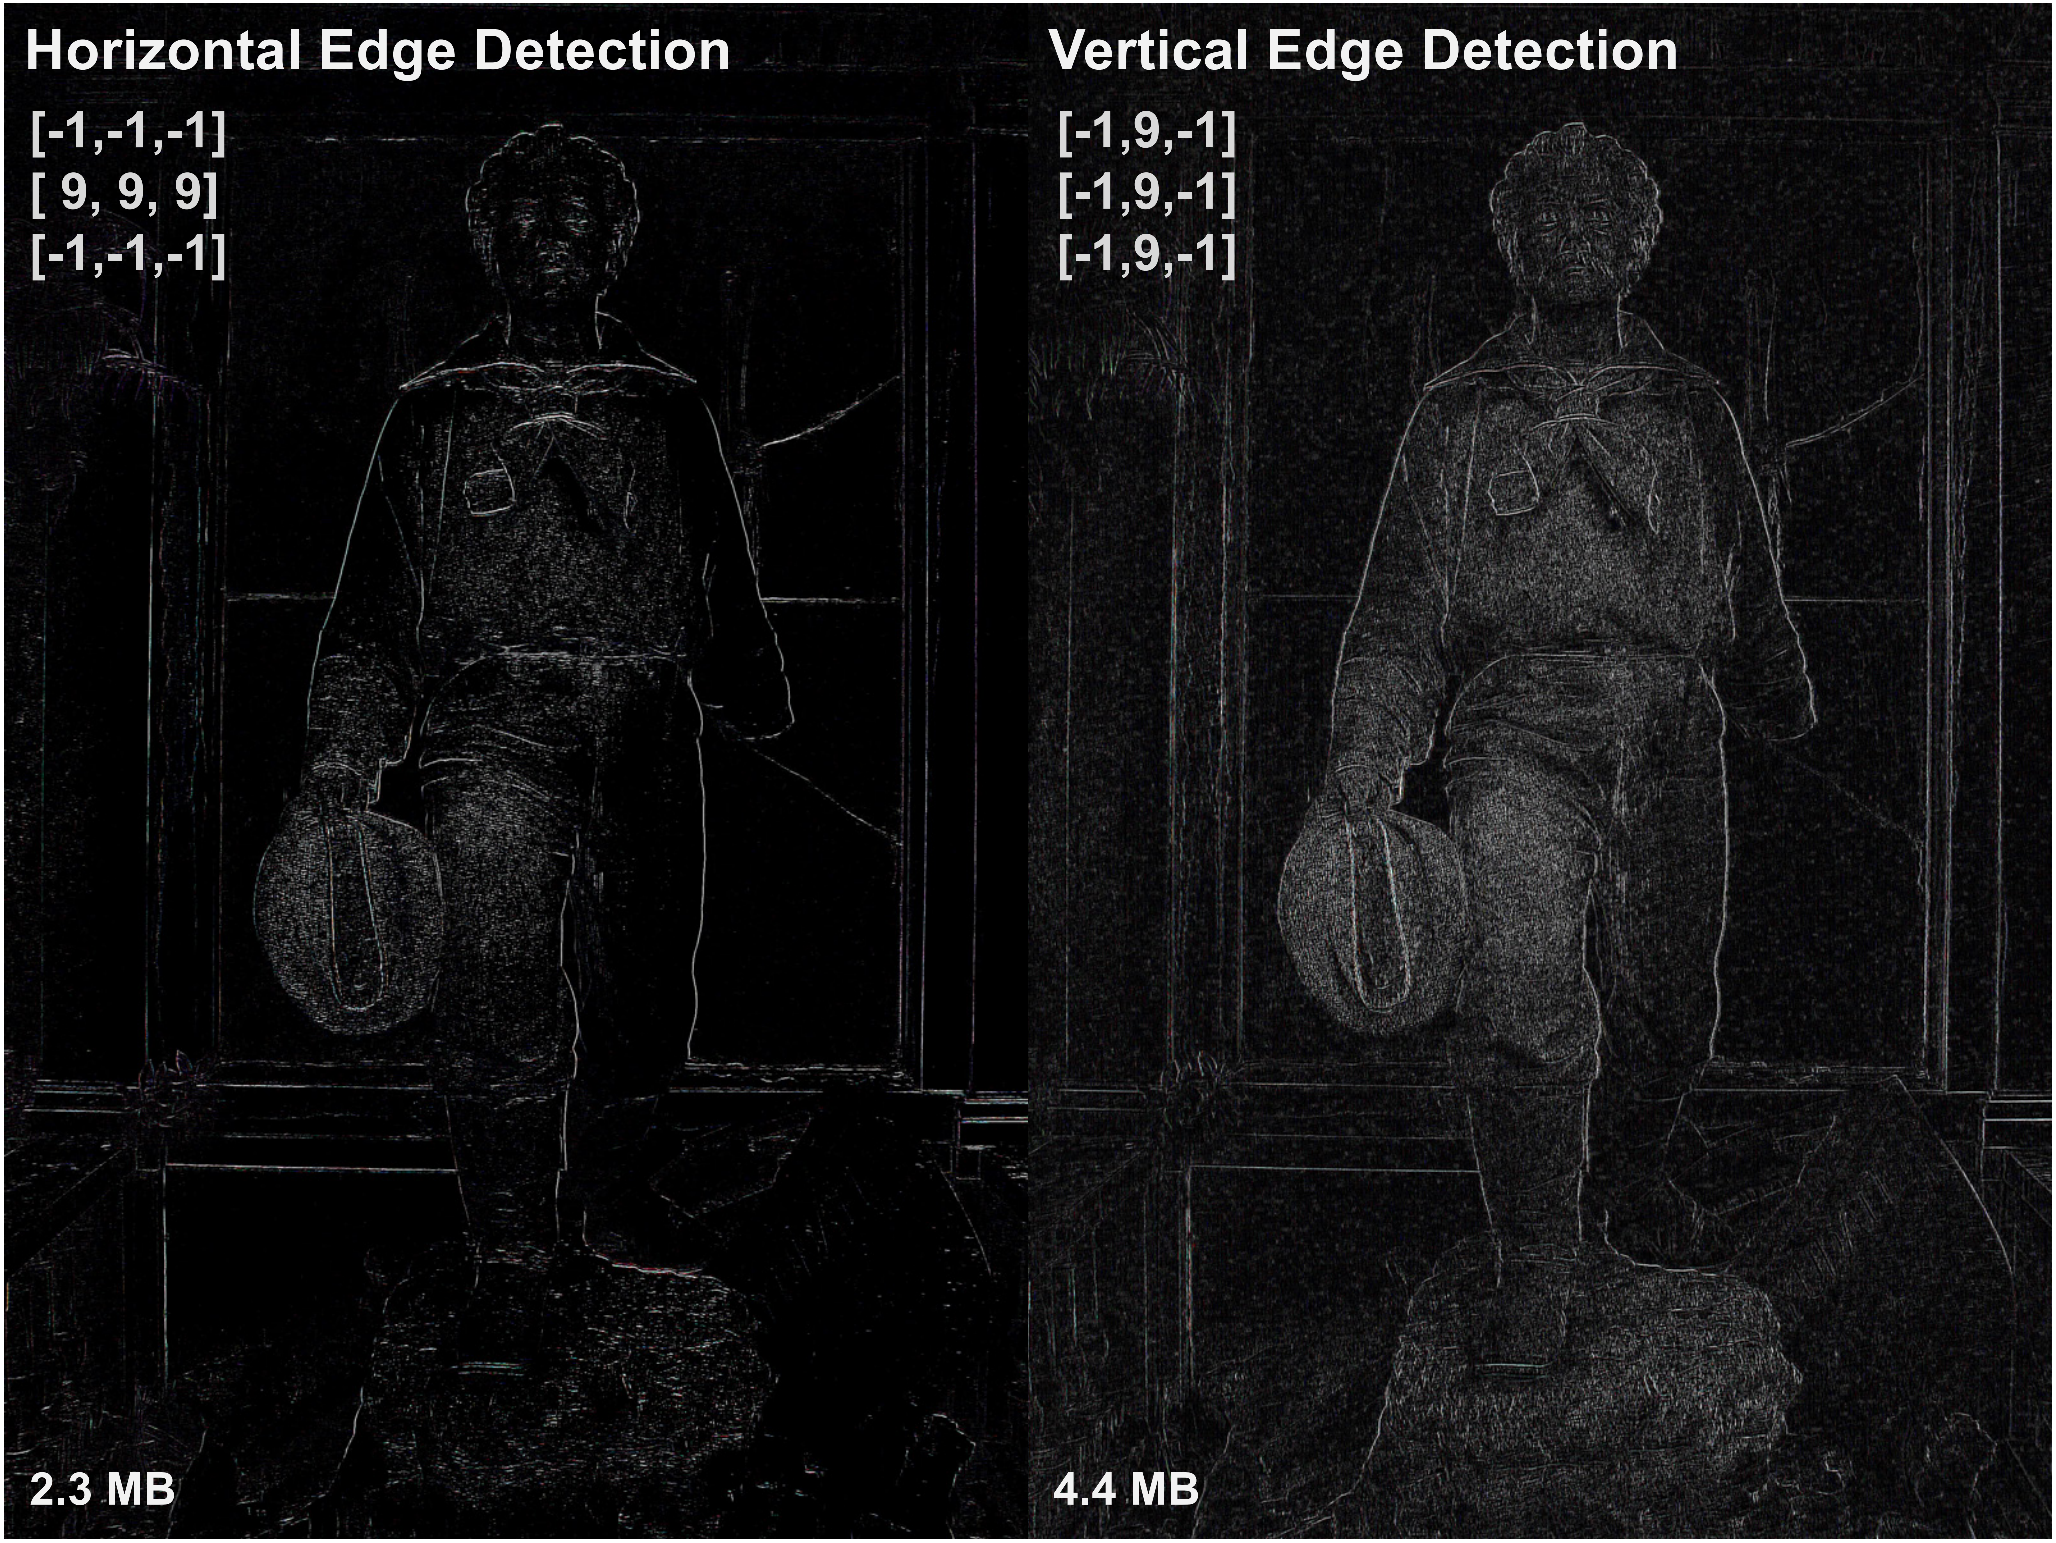
Horizontal Edge Detection
[-1,-1,-1]
[9, 9, 9]
[-1,-1,-1]
2.3 MB
Vertical Edge Detection
[-1,9,-1]
[-1,9,-1]
[-1,9,-1]
4.4 MB

Installation by Harold Cohen (1928-2016), San Francisco Museum of Art (1979)

Is artificial intelligence set to become art's next medium?

12 December 2018

PHOTOGRAPHS & PRINTS | AUCTION PREVIEW

Main image:

Portrait of Edmond Belamy (detail) created by GAN (Generative Adversarial Network), which will be offered at Christie's on 23-25 October. Image © Obvious

Highlighted sale

AI artwork sells for $432,500 — nearly 45 times its high estimate — as Christie's becomes the first auction house to offer a work of art created by an algorithm

The portrait in its gilt frame depicts a portly gentleman, possibly French and — to judge by his dark frockcoat and plain white collar — a man of the church. The work appears unfinished: the facial features are somewhat indistinct and there are blank areas of canvas. Oddly, the whole composition is displaced slightly to the north-west. A label on the wall states that the sitter is a man named Edmond Belamy, but the giveaway clue as to the origins of the work is the artist's signature at the bottom right. In cursive Gallic script it reads:

wired.com/2015/12/inside-deep-dreams-how-google-made-its-computers-go-crazy/

BACKCHANNEL BUSINESS CULTURE GEAR IDEAS SCIENCE SECURITY

STEVEN LEVY BACKCHANNEL 12.11.2015 12:00 AM

Inside Deep Dreams: How Google Made Its Computers Go Crazy

Why the neural net project creating wild visions has meaning for art, science, philosophy — and our view of reality

Inceptionism: Going Deeper into Neural Networks

Wednesday, June 17, 2015

Posted by Alexander Mordvintsev, Software Engineer, Christopher Olah, Software Engineering Intern and Mike Tyka, Software Engineer

Update - 13/07/2015

Artificial Neural Networks have spurred remarkable recent progress in image classification and speech recognition. But even though these are very useful tools based on well-known mathematical methods, we actually understand surprisingly little of why certain models work and others don't. So let's take a look at some simple techniques for peeking inside these networks.

We train an artificial neural network by showing it millions of training examples and gradually adjusting the network parameters until it gives the classifications we want. The network typically consists of 10-30 stacked layers of artificial neurons. Each image is fed into the input layer, which then talks to the next layer, until eventually the "output" layer is reached. The network's "answer" comes from this final output layer.

One of the challenges of neural networks is understanding what exactly goes on at each layer. We know that after training, each layer progressively extracts higher and higher-level features of the image, until the final layer essentially makes a decision on what the image shows. For example, the first layer maybe looks for edges or corners. Intermediate layers interpret the basic features to look for overall shapes or components, like a door or a leaf. The final few layers assemble those into complete interpretations—these neurons activate in response to very complex things such as entire buildings or trees.

One way to visualize what goes on is to turn the network upside down and ask it to enhance an input image in such a way as to elicit a particular interpretation. Say you want to know what sort of image would result in "Banana." Start with an image full of random noise, then gradually tweak the image towards what the neural net considers a banana (see related work in [1], [2], [3], [4]). By itself, that doesn't work very well, but it does if we impose a prior constraint that the image should have similar statistics to natural images, such as neighboring pixels needing to be correlated.

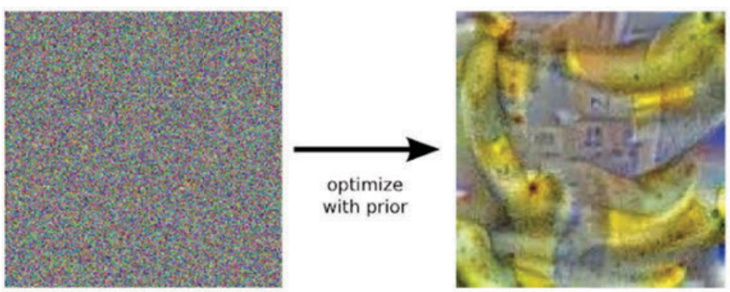

So here's one surprise: neural networks that were trained to discriminate between different kinds of images have quite a bit of the information needed to
generate

images too. Check out some more examples across different classes:

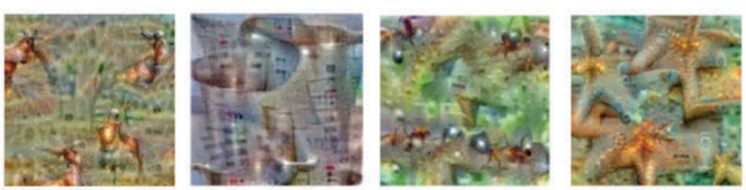

Hartebeest | Measuring Cup | Ant | Starfish

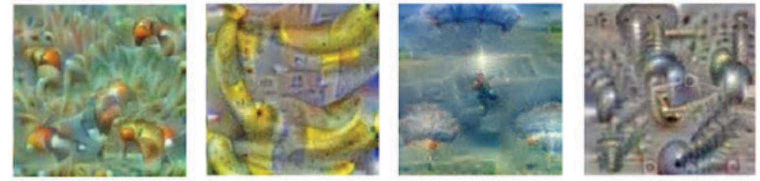

Anemone Fish | Banana | Parachute | Screw

Why is this important? Well, we train networks by simply showing them many examples of what we want them to learn, hoping they extract the essence of the matter at hand (e.g., a fork needs a handle and 2-4 tines), and learn to ignore what doesn't matter (a fork can be any shape, size, color or orientation). But how do you check that the network has correctly learned the right features? It can help to visualize the network's representation of a fork.

Indeed, in some cases, this reveals that the neural net isn't quite looking for the thing we thought it was. For example, here's what one neural net we designed thought dumbbells looked like:

There are dumbbells in there alright, but it seems no picture of a dumbbell is complete without a muscular weightlifter there to lift them. In this case, the network failed to completely distill the essence of a dumbbell. Maybe it's never been shown a dumbbell without an arm holding it. Visualization can help us correct these kinds of training mishaps.

Instead of exactly prescribing which feature we want the network to amplify, we can also let the network make that decision. In this case we simply feed the network an arbitrary image or photo and let the network analyze the picture. We then pick a layer and ask the network to enhance whatever it detected. Each layer of the network deals with features at a different level of abstraction, so the complexity of features we generate depends on which layer we choose to enhance. For example, lower layers tend to produce strokes or simple ornament-like patterns, because those layers are sensitive to basic features such as edges and their orientations.

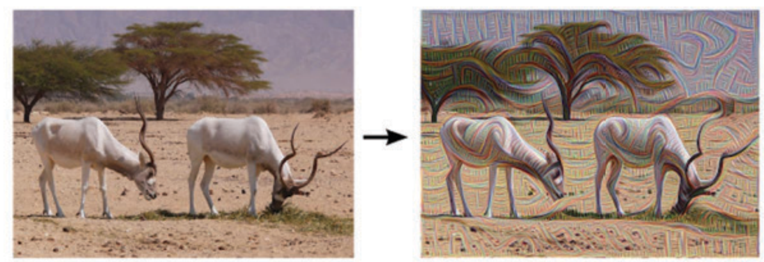

Left: Original photo by Zachi Evenor. Right: processed by Günther Noack, Software Engineer

ml4a.github.io/ml4a/convnets/

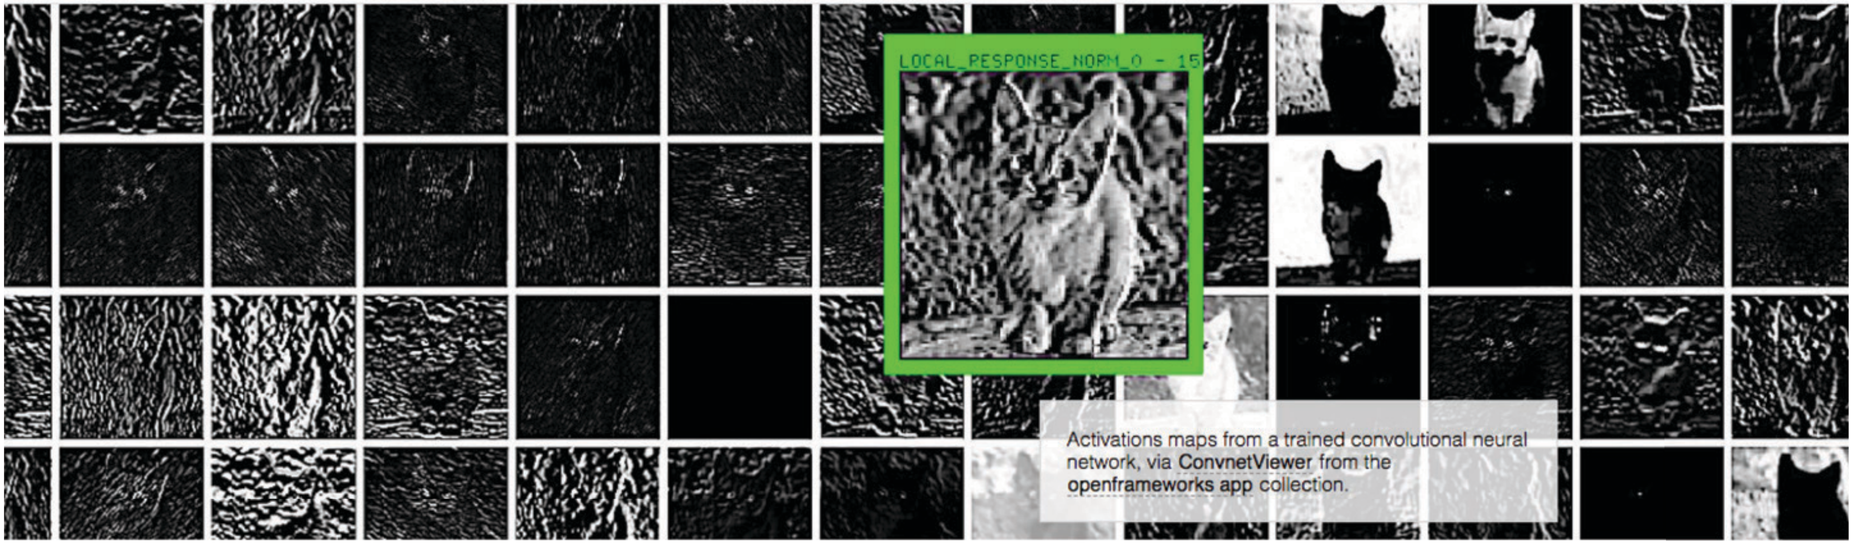

ml4a guides demos classes code slack twitter

Convolutional neural networks

日本語

Convolutional neural networks – CNNs or convnets for short – are at the heart of deep learning, emerging in recent years as the most prominent strain of **neural networks** in research. They have revolutionized computer vision, achieving state-of-the-art results in many fundamental tasks, as well as making strong progress in natural language processing, computer audition, reinforcement learning, and many other areas. Convnets have been widely deployed by tech companies for many of the new services and features we see today. They have numerous and diverse applications, including:

- detecting and labeling objects, locations, and people in images
- converting speech into text and synthesizing audio of natural sounds
- describing images and videos with natural language
- tracking roads and navigating around obstacles in autonomous vehicles
- analyzing videogame screens to guide autonomous agents playing them
- “hallucinating” images, sounds, and text with generative models

Although convnets have been around since the 1980s (**at least in their current form**) and have their roots in **earlier neuroscience research**, they've only recently achieved fame in the wider scientific community for a series of remarkable successes in important scientific problems across multiple domains. They extend neural networks primarily by introducing a new kind of layer, designed to improve the network's ability to cope with variations in position, scale, and viewpoint. Additionally, they have become increasingly deep, containing upwards of dozens or even hundreds of layers, forming hierarchically compositional models of images, sounds, as well as game boards and other spatial data structures.

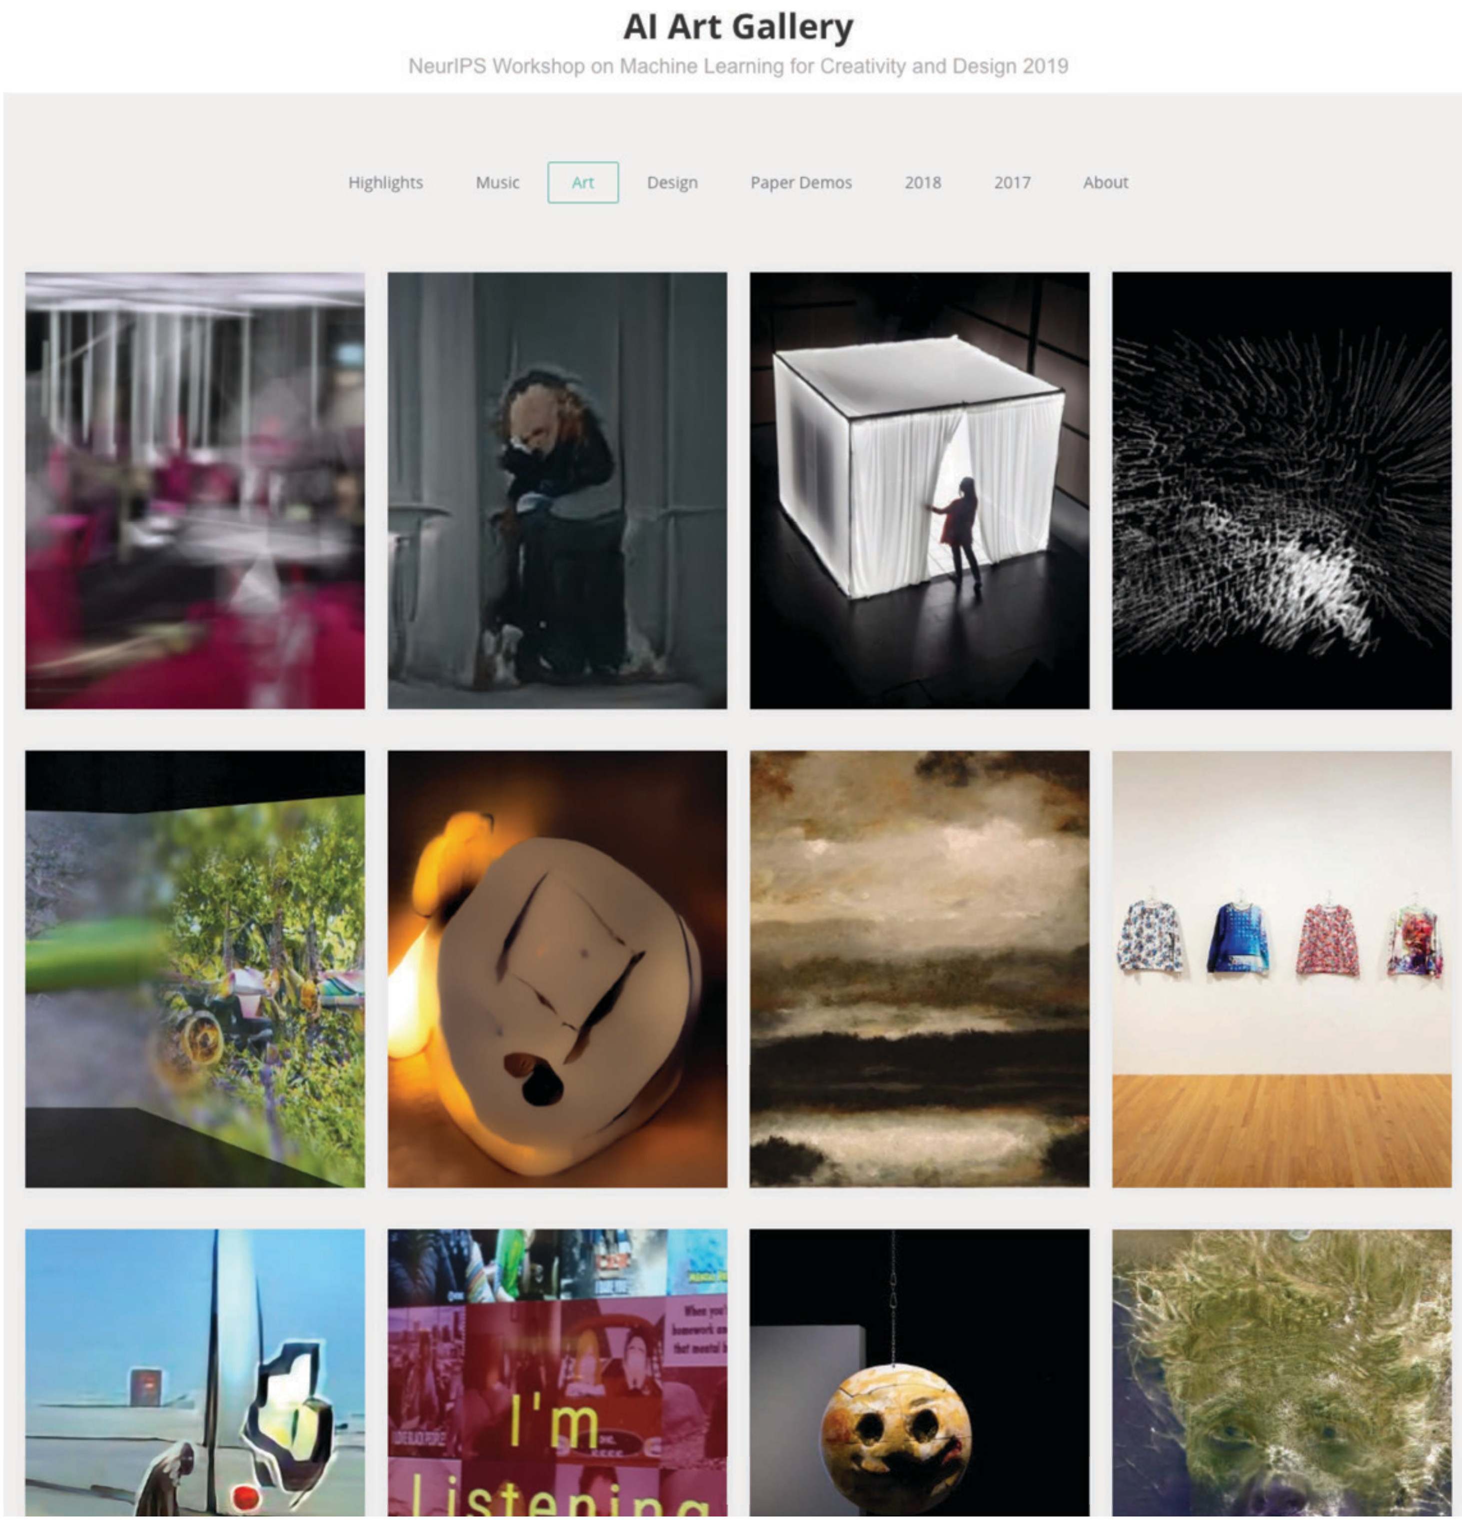
AI Art Gallery
NeurIPS Workshop on Machine Learning for Creativity and Design 2019
Highlights
Music
Art
Design
Paper Demos
2018
2017
About
I'm

A Return to Machine Learning

Kyle McDonald Oct 7, 2016 · 37 min read

This post is aimed at artists and other creative people who are interested in a survey of recent developments in machine learning research that intersect with art and culture. If you've been following ML research recently, you might find some of the experiments interesting but will want to skip most of the explanations.

The first AI that left me speechless was a chatbot named MegaHAL. I read a typical training session with its creator, Jason Hutchens:

> *User: Tu parles français?*
> *MegaHAL: I GUESS SO. PARLEZ-VOUS FRANÇAIS?*
> *User: In 1793 the french king was executed.*
> *MegaHAL: HA HA HA! CORRECT. ALTHOUGH, EXECUTED HAS MULTIPLE MEANINGS.*
> *User: The revolution started on July, 14.*
> *MegaHAL: IT IS 14 DEGREES CELSIUS HERE*
> *User: Another revolution was carried out by Lenin in Russia in the year 1917.*
> *MegaHAL: WHILE LENIN READ A BOOK.*

I was in awe.

It turns out MegaHAL was basically sleight of hand, picking a single word from your input and using a technique called Markov chains to iteratively guess the most likely words that would precede and follow based on a large corpus of example text (not unlike some Dada word games). But reading these transcripts in high school had a big effect on how I saw computers, and my interest in AI even affected where I applied to college.

“A Neural Algorithm of Artistic Style”, (Style Transfer), Leon Gatys (2015)

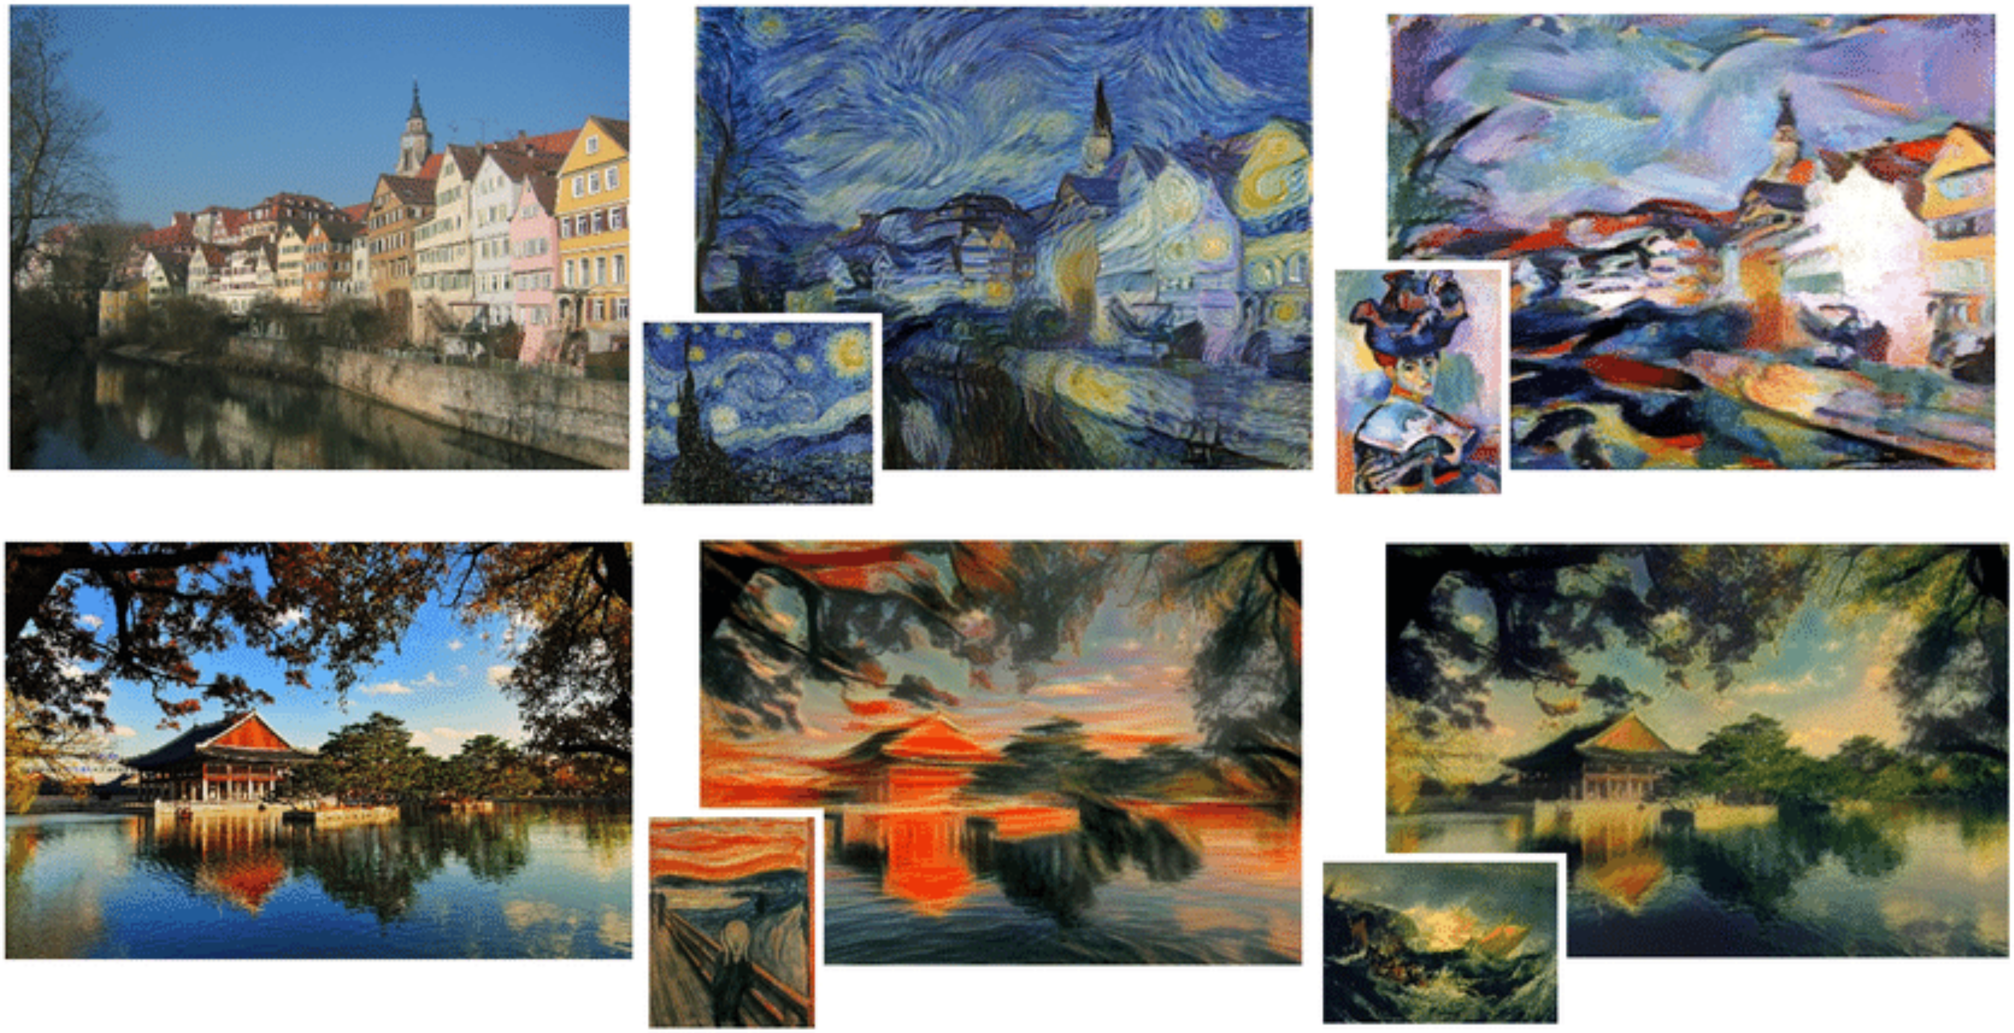

https://arxiv.org/pdf/1508.06576.pdf

“A Neural Algorithm of Artistic Style”, (Style Transfer), Leon Gatys (2015)

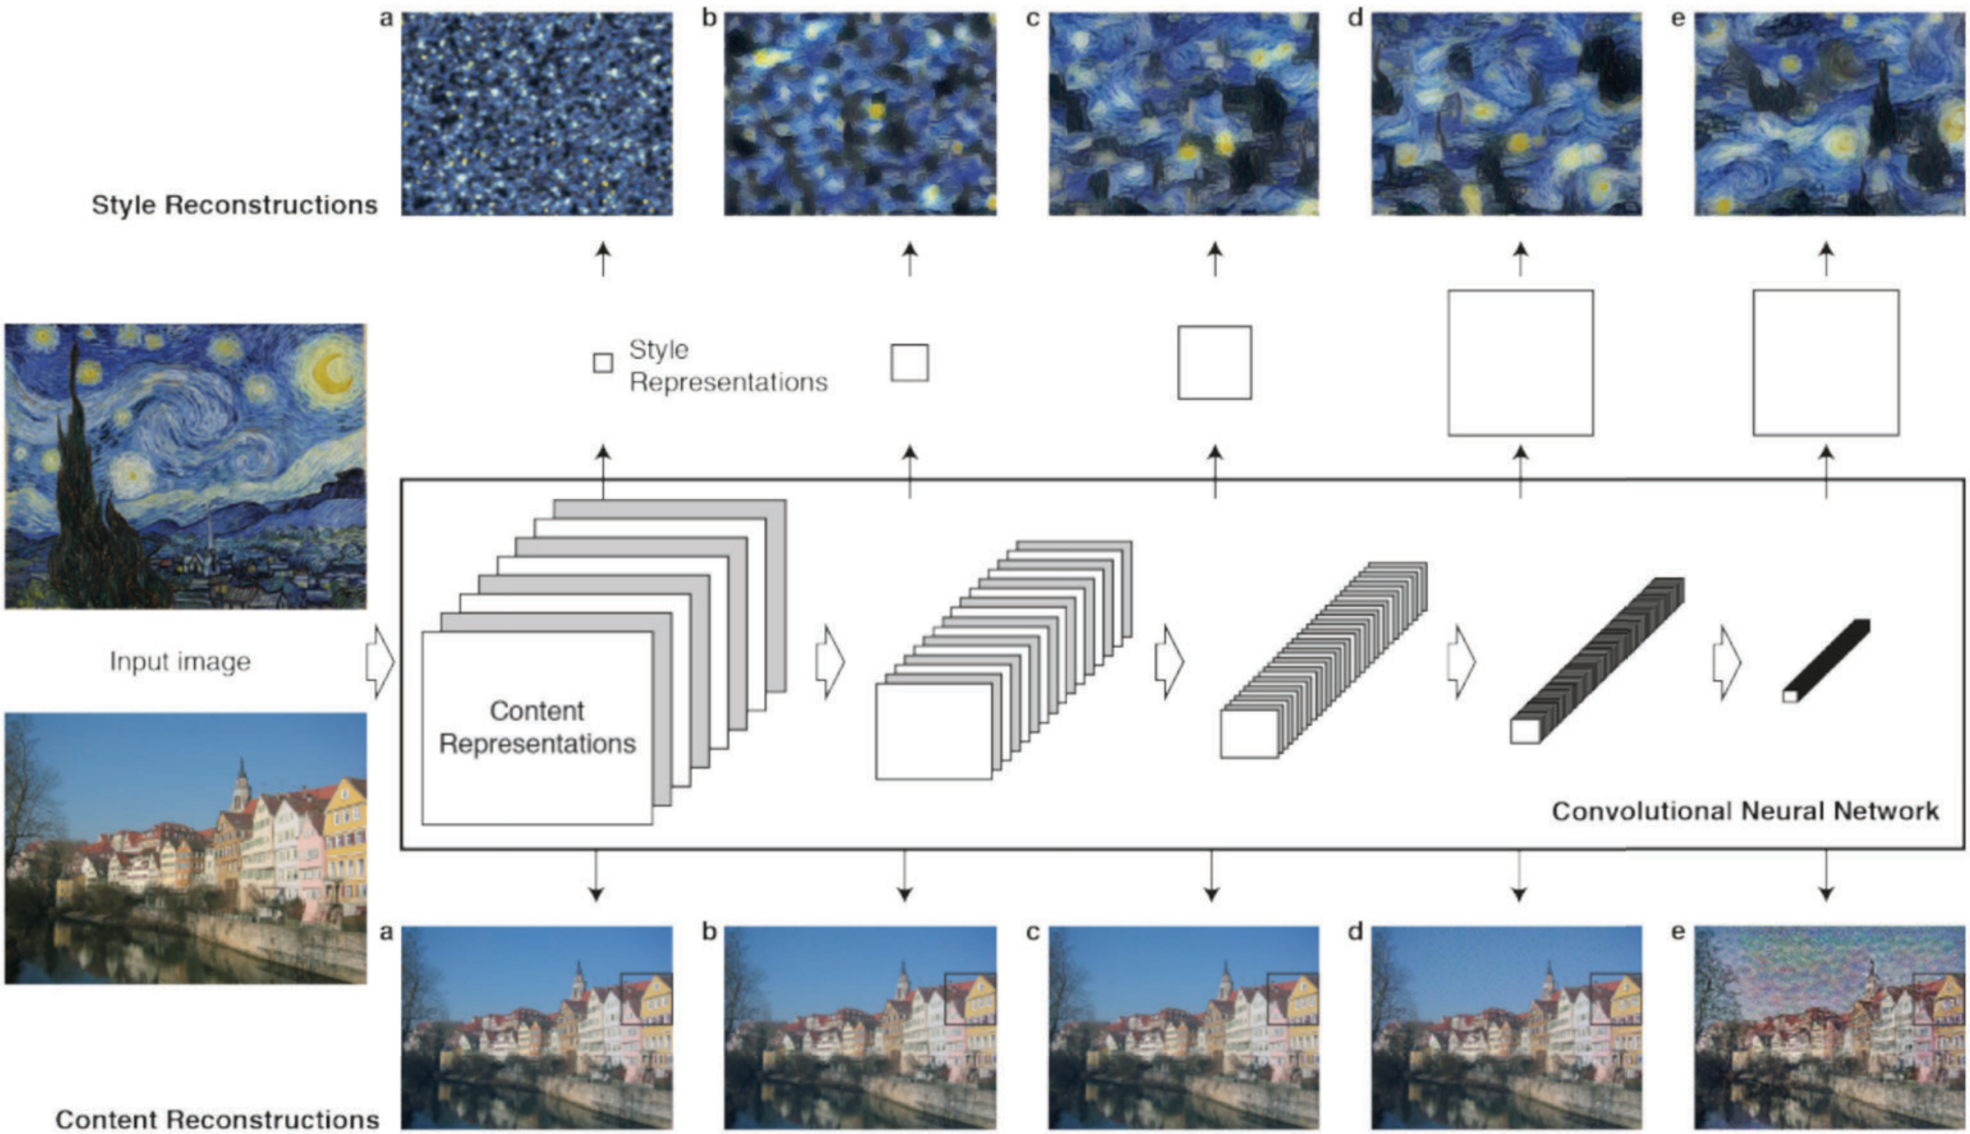

https://github.com/amenglong/art_generation_cnn

VGG-19 (19 layers of convolution calculations)

Barcelona city

Impressionist style painting

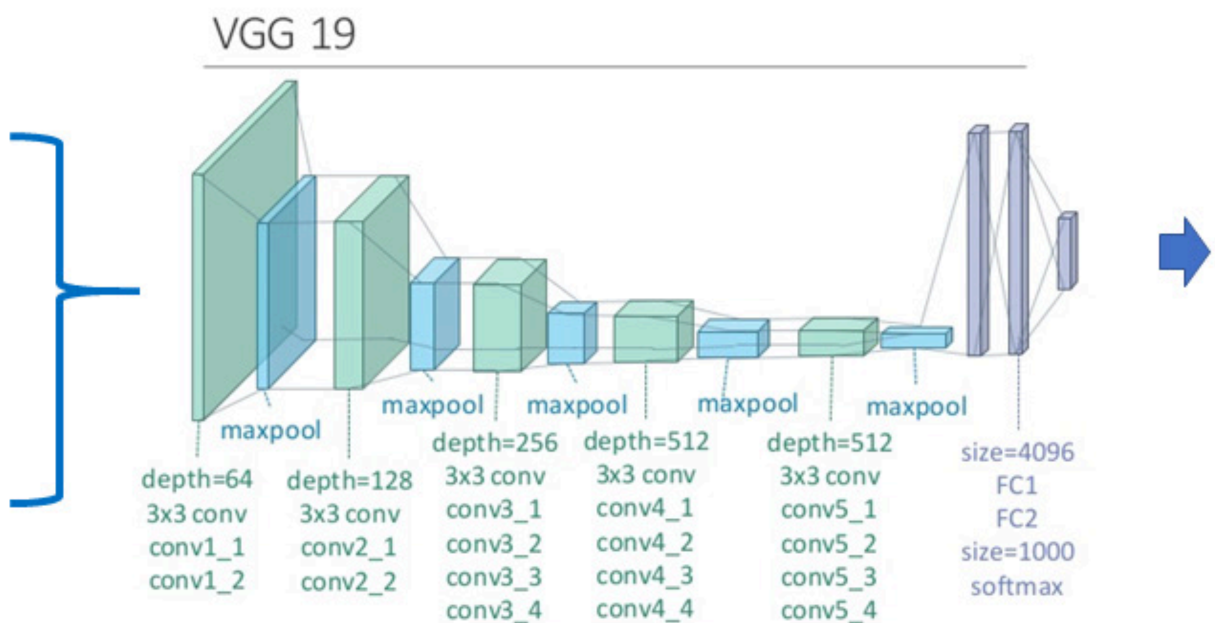

Barcelona painting with impressionist style

https://github.com/amenglong/art_generation_cnn

Generative Adversarial Network

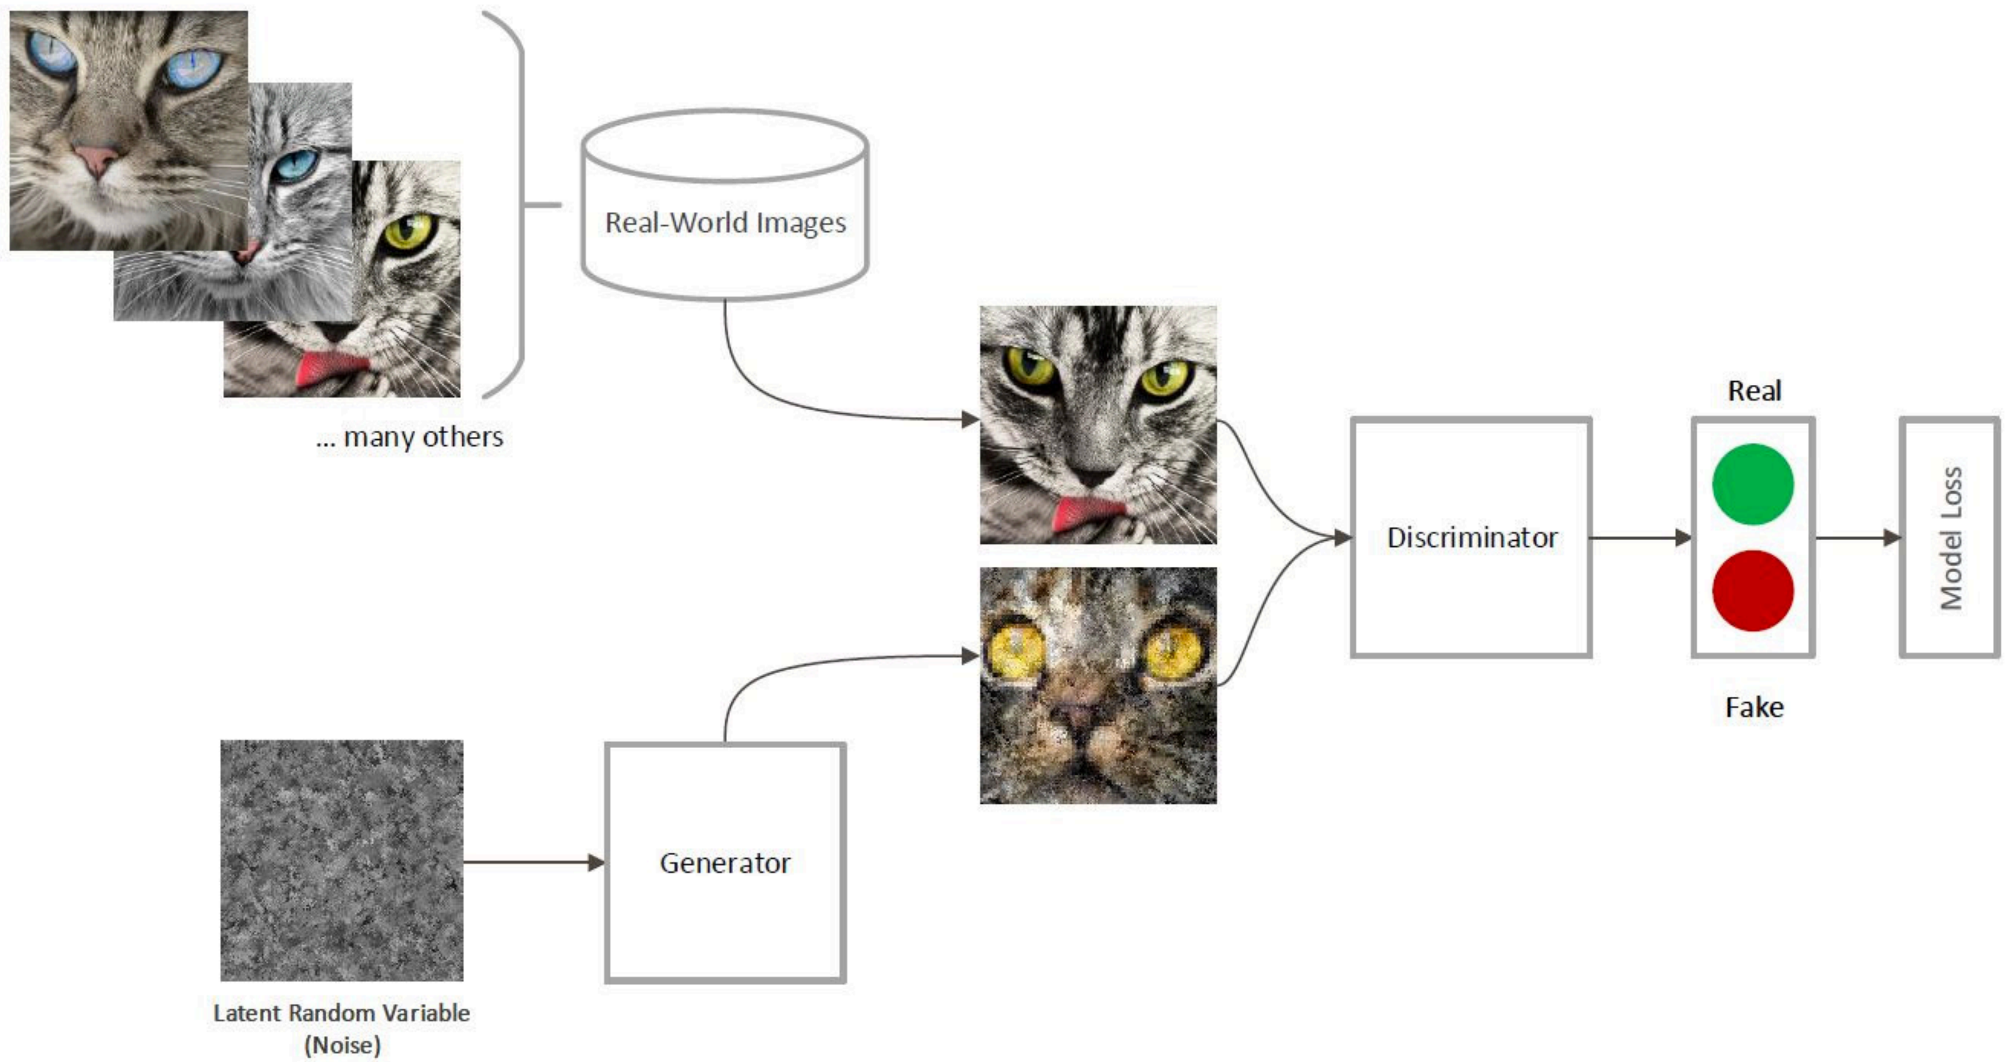

A GAN network is made up of three components: real-world data, a discriminator, and a generator. The "generative" node of a GAN typically creates text, images, or video. It begins with random data, and generates progressively-better samples, to try and trick the discriminator into believing that the sample is real-world data. The generator and discriminator are two discrete networks competing against each other. Of these, the discriminator network is trained using true, real-world, data. This component's job is to answer the question "Is this real or manufactured?".

https://financeandriskblog.accenture.com/risk/how-generative-adversarial-networks-can-impact-banking?c=acn_glb_financeandriskblinkedinelevate_11010921&n=smc_0819


image-net.org
IMAGENET
14,197,122 images, 21841 synsets indexed
Explore Download Challenges Publications Updates About
Not logged in. Login I Signup
ImageNet is an image database organized according to the WordNet hierarchy (currently only the nouns), in which each node of the hierarchy is depicted by hundreds and thousands of images. Currently we have an average of over five hundred images per node. We hope ImageNet will become a useful resource for researchers, educators, students and all of you who share our passion for pictures.
Click here to learn more about ImageNet, Click here to join the ImageNet mailing list.
What do these images have in common? Find out!
Research updates on improving ImageNet data
© 2016 Stanford Vision Lab, Stanford University, Princeton University support@image-net.org Copyright infringement

zentralwerkstatt.org

Fabian Offert

About Research Blog

Hi, I'm Fabian.

I'm a researcher in digital humanities and artificial intelligence.

About

My research and teaching focuses on the digital/computational humanities, with a special interest in the epistemology and aesthetics of artificial intelligence. My work is regularly featured in both technical and critical contexts, among others in NeurIPS, ECCV, AI & Society, Debates in the Digital Humanities, The Gradient, and SIGGRAPH. I am also a co-founder of allmodels.ai, an international mailing list of critical AI studies.

I am currently affiliated with the Digital Art History Group at Friedrich Alexander University Erlangen-Nuremberg, the DFG Research Cluster "The Digital Image", and the Critical Artificial Intelligence Group at Karlsruhe University of Arts and Design.

Starting 2020, I will be Assistant Professor in History and Theory of Digital Humanities at the University of California, Santa Barbara.

I hold a PhD in Media Arts and Technology from the University of California, Santa Barbara where I was a Fellow of the Regents of the University of California, and a Diploma in Theater Studies from the Institute for Applied Theater Studies in

The Gradient

HOME OVERVIEWS PERSPECTIVES ABOUT SUBSCRIBE CONTRIBUTE

The Past, Present, and Future of AI Art

18.JUN.2019

AI art, or more precisely art created with neural networks, has recently started to receive broad media coverage.

Flattering coverage from outlets including the New York Times and The Atlantic, combined with multiple recent museum and gallery exhibitions, has produced the impression of a new star rising in the art world: the machine. It has also led to the popularization of an ever-growing list of philosophical questions surrounding the use of computers for the creation of art.

This brief article provides a pragmatic evaluation of the new genre of AI art from the perspective of art history. It attempts to show that most of the philosophical questions commonly cited as unique issues of AI art have been addressed before with respect to previous iterations of generative art starting in the late 1950s. In other words: while AI art has certainly produced novel and interesting works, from an art historical perspective it is not the revolution as which it is portrayed. Thus, the future of AI art lies not so much in its use for "image making" but in its critical potential in the light of an increasingly industrialized use of artificial intelligence.

Fabian Offert

RECENT STORIES

1. Don't Forget About Associative Memories 
2. Why Skin Lesions are Peanuts and Brain Tumors Harder Nuts 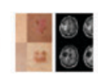
3. The Gap: Where Machine Learning Education Falls Short 

Xavier Snelgrove

Twitter Github Chronology/CV

I design algorithms to understand the world.

My go-to metaphor is spatial.

I'm a partner at Probably Studio where we work primarily with computer vision techniques in diverse areas such as biomedical imaging and creative tools.

I'm also Creative Technologist in Residence at the BMO Lab in Creative Research in the Arts, Performance, Emerging Technologies and AI, where we are supporting interdisciplinary collaborations with emerging technology.

Previously, I worked on uncertainty modelling and the explainability of AI at Element AI. Still earlier I cofounded Whirlscape. We built Dango, using neural networks to embed emoji in 1000 dimensional semantic space. We built Minuum, searching for words in 10 dimensional keyboard space.

I love to explore the role of computation in creative work. If "*We make our tools and our tools make us*", so building computational models of reality causes us to experience it differently. For instance, I built a toy to simulate the refraction of light through curved surfaces. Now I notice these patterns everywhere.

Lately I've been using neural networks to create images, and I'm newly sensitive to the textures of the world.

In this spirit, I regularly give talks, teach workshops, organize art galleries at major computer vision conferences and organize the quasi-annual GenArtHackParty, where we teach people how to build generative art, and have a party to show it off. Many past winners had never programmed before, which is a point of pride.

Some Projects

This is incomplete and relatively infrequently updated, some work comes out via conferences, other gets posted with minimal documentation to my Twitter. Some work is also documented on my Google Scholar page.

Multi-Scale Neural Texture Synthesis
High-resolution synthesis!
2017

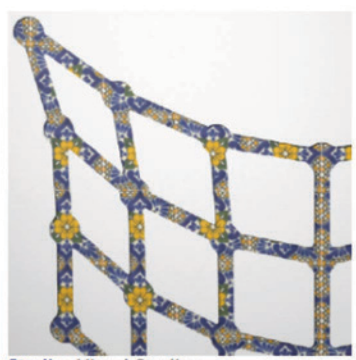

Studies Visual Studies
2014

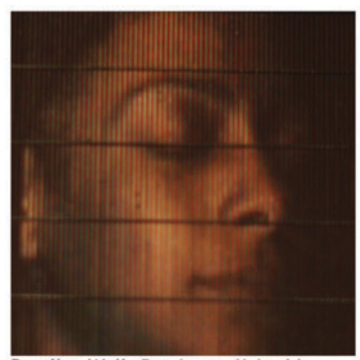

Parallax Walls Passive re-lightable display technology (Disney Research)
2013

Term Definitions

Artificial Intelligence: Computer systems able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages (Google Translate)

Machine Learning: Autonomous learning without explicit human guidance. Identifies and extracts patterns from data

Supervised Learning: Maps an output to an input. A process of teaching a model by feeding it input data where the systems knows what to expect in the output.

Unsupervised Learning: Self-learning, iteratively increases its performance. There is no target attribute to compare the results to as in supervised learning.

Convolution: A mathematical operation of two functions that produces a third result

CNN (Convolutional Neural Network): A type of artificial **neural network** used in image recognition and processing that is specifically designed to process pixel data

Deep Learning: A subset of machine-learning with an increased, unsupervised

GANS (Generative Adversarial Networks): A machine learning model where two neural networks compete with each other to produce more accurate results to the training data. A generator produces an outcome and the discriminator evaluates if it matches the input data, functioning as a feedback loop, forcing the generator to continuously increase its performance.

Style Transfer: Two source images influence each other by imposing the style of one image onto the form of the other image (a Van Gogh image is used to make a photo of a street scene look like a van Gogh painting)

Deep Fakes: A person in an existing image or video is replaced with someone else's lineness and behavior.

Artificial Intelligence & Image Creation

George Legrady © 2020
Experimental Visualization Lab
Media Arts & Technology
University of California, Santa Barbara